



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

IN RE APPLICATION OF:
RAVI KUMAR ARIMILLI, ET AL.

SERIAL NO.: **09/753,052**

FILED: **DECEMBER 28, 2000**

FOR: **SUPPORTING ADDITION OF
HETEROGENOUS
PROCESSORS IN A
SYMMETRIC
MULTIPROCESSING (SMP)
SYSTEM**

ATTY. DOCKET NO.: **AUS920000679US1**

§

§ EXAMINER: **CHARLES HARKNESS**

§

§

§

§ ART UNIT: **2183**

§

§

§

§

§

§

§

APPEAL BRIEF UNDER 37 C.F.R. §1.192

Mail Stop Appeal Briefs - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

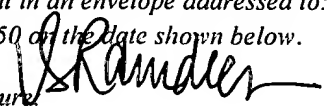
Sir:

This Brief is submitted in support of the Appeal of the Examiner's final rejection of Claims 1-20 in the above-identified application. A Notice of Appeal was filed in this case on November 19, 2004 and received in the United States Patent and Trademark Office on November 19, 2004. A one month extension of time is required and is requested. A check in the amount of \$120.00 for the extension of time is attached. Please charge the fee of \$500.00 due under 37 C.F.R. §1.17(c) for filing the brief, as well as any additional required fees, to **IBM Deposit Account No. 09-0447**.

Certificate of Transmission/Mailing

*I hereby certify that this correspondence is being facsimile transmitted to the USPTO at 703-872-9306 or deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:
Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450 on the date shown below.*

Typed or Printed Name: Shenise Ramdeen Date: February 21, 2005

Signature: 

AUS920000679US1

Appeal Brief

Serial No. 09/753,052



REAL PARTY IN INTEREST

The real party in interest in the present Application is International Business Machines Corporation, the Assignee of the present application as evidenced by the Assignment set forth at reel 011428, frame 0631.

RELATED APPEALS AND INTERFERENCES

There are no other appeals or interferences known to Appellants, the Appellants' legal representative, or assignee, which directly affect or would be directly affected by or have a bearing on the Board's decision in the pending appeal.

STATUS OF CLAIMS

Claims 1-20 stand finally rejected by the Examiner as noted in the Final Office Action dated October 7, 2004. The rejection of Claims 1-20 is appealed.

STATUS OF AMENDMENTS

Following the Office Action dated March 18, 2004, Appellants submitted an amendment on June 18, 2004, which was entered by the Examiner. No amendments have been made subsequent to the Final Office Action.

SUMMARY OF THE CLAIMED SUBJECT MATTER

Appellants' invention provides a multiprocessor data processing system designed with hardware and software components that support a later addition of heterogeneous processors, where each heterogeneous processor has unique operating characteristics including, for example, different processing speeds (frequency), different integrated circuit design, different cache topologies (sizes, levels, etc.). The data processing system includes a specialized set of pins associated with the system bus or switch for connecting the heterogeneous processors. The data processing system also includes an enhanced operating system (OS) with enhanced communication protocol that supports the addition of heterogeneous processors and enables sharing of the workload among the various interconnected processors, each operating (processing data, etc.) according to its individual configuration (c.g., at respective frequencies).

Second and third generation heterogeneous processors are connected to the specialized set of pins, which allow the newer added processors to support enhanced system bus protocols with downward compatibility to the previous generation processors. The enhanced OS, communication protocol, and other inter-processor logic enable the heterogeneous multiprocessor data processing system to process workload in an efficient manner similarly to a symmetric multiprocessor system.

As recited by the independent claims, namely exemplary Claim 1, the invention provides the following key features:

interconnection means for later connecting a second processor ... heterogeneous to said first processor, said interconnection means enables ... collectively operate as a symmetric multiprocessor (SMP) system; and an enhanced operating system (OS) that supports ... cache coherency operations based on a collective memory configuration of the SMP, wherein said OS logs operating characteristics and cache topology data of each processor ... to calculate a most efficient work allocation among processors; wherein said interconnection means and said enhanced operating system support backward and forward compatibility ... inter-processor operations including cache intervention, prefetching, and intelligent cache states.
(emphases added)

Other features provided by Appellants' claims include (among others):

- (1) "interconnect means ...comprising system data bus, base address bus, **master processor select bus**, base snoop response bus and **extended snoop response bus**, wherein each bus includes one or more pins that are set/reset to indicate a particular condition of a connected component" (Claim 4, *emphases added*);
- (2) "wherein said **master processor select bus** includes ... pins, each connected to an added processor, wherein when one of said pins is set to an active state, the connected processor operates as a master" (Claim 5, *emphasis added*);
- (3) "pin is set when a read operation is issued.. master processor; and snooped by a second added processor with cache line in the R coherency state, ... drives the extended snoop response bus with shared intervention information..." (Claim 6, *emphases added*).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

- A. The Examiner's rejection of Claims 1-20 under 35 U.S.C. §103(a) as being unpatentable over *McCrory* (U.S. Patent No. 6,513,057) in view of *Jayakumar* (U.S. Patent No. 5,904,733) and further in view of *Derrick, et al.* (U.S. Patent No. 5,704,058) is to be reviewed on Appeal;

ARGUMENT

- A. The rejection of Claims 1-20 under 35 U.S.C. §103(a) as being unpatentable over *McCrory* (U.S. Patent No. 6,513,057) in view of *Jayakumar* (U.S. Patent No. 5,904,733) and further in view of *Derrick, et al.* (U.S. Patent No. 5,704,058) is not well founded and should be reversed.

1. No Motivation to Combine References

First, there is no motivation in either reference for combining the above references. In fact, absent the teachings of Appellants' claimed invention, the combination would probably not have been considered by the Examiner. While heterogeneous processing systems were known, all such systems were static, i.e., prefabricated and provided specific supporting communication protocols and inter-operability hardware/software logic, etc. Prior to Appellants' invention, there was no consideration given to the possibility of enabling a later add-on of a heterogeneous processor to an existing system. Conventional thought at that time was that such add-ons would lead to instability in the main system and may even lead to the system crashing since the OS, communication protocols, cache configuration and coherency protocols, etc. required for the heterogeneous processor would be too different from those existing in the main system to support the later add-on of the heterogeneous processor.

McCrory describes a heterogeneous multiprocessor with "processors from distinct families ...**integrated on a single platform**" and "coupled with an **implementation specific communication mechanism** through **family specific bus interface converters**" (Abstract; *emphases added*). The processors in *McCrory's* system are specifically described as "integrated" or "packaged". Neither of these two terms connotes the functionality of being able

to later add a heterogeneous processor via an interconnect means that is designed specifically for that functionality.

Jayakumar, in stark contrast, specifically describes adding **homogenous processors** to a **traditional SMP** environment (i.e., one with homogenous processors only). *Jayakumar* briefly mentions (in the background) that “additional processors may be added to the SMP without alerting the software.” It is, however, very clear to one skilled in the art at the time of Appellants’ invention that *Jayakumar* discussion covered only homogenous systems.

Of significance here is that *Jayakumar*’s patent is restricted to homogenous add-ons while *McCrory*’s patent is specifically limited to an integrated heterogeneous environment on a single, non-expandable platform with “implementation specific” software/hardware support. Both references conform to what was known in the art prior to Appellants’ invention. Examiner admits that *Jayakumar* teaches “a traditional SMP environment, while *McCrory* has taught a heterogeneous SMP system.”

Appellants acknowledge, in the applications’ background section, that addition of homogenous processors to an SMP was known in the art. Extend this add-on functionality (limited to homogenous processors) to a processor environment that would support later add-on of heterogeneous processors, running at different frequencies and having different cache configurations, etc., was, however, not previously available or contemplated.

As noted above, Examiner does recognize the inherent differences in the homogenous environment and the heterogeneous environment described by the references. It is clear from the limitations built into the respective references that absent the teachings of Appellants’ invention, one skilled in the art would not contemplate a combination of *McCrory* with *Jayakumar*. Since Examiner may not rely on the teachings of Appellants’ invention and use hindsight reasoning to support a combination that would not have otherwise been made, the above combination cannot be utilized to support the rejection of Appellants’ claims. The rejection of Appellants’ Claims 1-20 is therefore not well founded and should be reversed.

2. The Combination Does Not Suggest Appellants' Claimed Invention

Even in one would be inclined to combine the references, the combination still does not suggest the novel features provided by Appellants' claimed invention, as described below with reference to specific claims.

i. Claims 1-3 10, 12, 14-15 "interconnect means and OS functionality"

Appellants' exemplary Claim 1 recites:

*interconnection means for later connecting a second processor ... heterogeneous to said first processor...; and
an enhanced operating system (OS) that supports inter-processing operations ... including cache coherency operations based on a collective memory configuration of the SMP, wherein said OS logs operating characteristics and cache topology data of each processor connected to the interconnection means to calculate a most efficient work allocation among processors;*
(emphases added).

On page three of the Action, Examiner admits that *McCrory* does not teach or suggest the "interconnection means for later connecting a second [heterogeneous] processor." In fact, all processors in *McCrory's* system are specifically described as "integrated" or "packaged" on a single chip/backplane. Examiner references Col. 1, lines 10-26 of *Jayakumar* to support the rejection of this feature. That section of *Jayakumar* states that "*additional processors may be added to the SMP without alerting the software*" (emphasis added).

Examiner rationalizes that despite the fact that *Jayakumar* teaches adding homogenous processors to a traditional SMP environment and *McCrory* teaches an integrated heterogeneous SMP system, one skilled in the art "would recognize the benefit it (*sic*) adding additional processors, from different families, ...," and, therefore, it would have been obvious "to have interconnection means for later connecting a second, heterogeneous processor to increase the system's processing power."

What Examiner fails to include in his analysis is a basic understanding that a later addition of a homogenous processor (which does not require "alerting the software") is not synonymous with and on several orders of magnitude easier than providing a computer system

architecture that supports a later addition of a heterogeneous processor. The enhancements required in the software and the necessity of alerting the software (e.g., the OS) whenever a heterogeneous processor was connected to the system in order for the OS to dynamically provide support for a resulting heterogeneous multi-processor system (all post-manufacture of the main computer system) is not suggested or contemplated by the combination of references. These functional elements are unique to Appellants' invention.

As mentioned in section 1 above, each of the patents being relied on is instructive in its limitations. The processor configuration of *McCrory* is static in nature, i.e., not capable of being changed/upgraded post-manufacture. The design and operational qualities for that specific set of processors are built into the system and cannot be changed at a later time. Thus *McCrory* clearly does not support a latter addition of another heterogeneous processor as *McCrory* would not be able to support the additional processor. *Jayakumar* would not be able to handle addition of a heterogeneous processor because of the lack of an enhanced OS required to detect and provide support for the processor differences without stalling or crashing the entire system and without sending all workload to one or the other of the processor in an unbalanced allocation.

Clearly, *McCrory* never suggests or contemplates supporting a later addition of a processor of any kind, and *Jayakumar* never suggests or contemplates that the SMP would support any addition of a heterogeneous processor and thus never contemplates the enhanced OS functionality required to support such a heterogeneous configuration.

With respect to this last feature, Examiner admits that the combination of *McCrory* and *Jayakumar* does not teach the enhanced OS required to support the inter-processing operations when a later add on of a heterogeneous processor occurs. Examiner relies on *Derrick* for support of the rejection of this element and associated features. Notably, Examiner focuses his arguments solely on *Derrick's* discussion of cache coherency operations and an arbitration scheme efficient allocation of cache bus bandwidth. *Derrick*, however, never addresses or suggests an enhance OS and/or certain other functions attributed to the enhanced Os such as providing a "most efficient work allocation among processors."

The cited sections of *Derrick*, namely, col. 1, lines 15-57 and col. 3, line 7-17, describes memory bandwidth. *Derrick*'s description of providing optimized bus bandwidth for the caches is not suggestive of the feature *Derrick* is provided to reject, namely an enhanced OS that supports later addition of heterogeneous processors. Further, providing optimized bus bandwidth is not synonymous with efficient workload allocation among multiple heterogeneous processors. It thus appears Examiner has mischaracterized what is actually taught by *Derrick* and that *Derrick* is devoid of several of the features Examiner attributed to *Derrick* in supporting the rejection of Appellants' claims.

For the above reasons, Examiner's rejection of Claims 1-3 10, 12, and 14-15 is not well founded and should be reversed.

ii. Claims 4-5, 11, and 17-18 (interconnect pins and specialized busses)

Appellants' Claim 4 further recites:

*interconnect means ...comprising system data bus, base address bus, **master processor select bus**, base snoop response bus and **extended snoop response bus**, wherein each bus includes one or more pins that are set/reset to indicate a particular condition of a connected component*

(*emphasis added*); and Claim 5 recites:

*wherein said master processor select bus includes ... **pins**, each connected to an added processor, **wherein when one of said pins is set to an active state, the connected processor operates as a master***

(*emphasis added*).

None of the above features are taught or suggested by any of the references or the combination of references. There is also no suggestion of connecting the new heterogeneous processors to the existing system via specialized connection pins and/or special buses, which include a master processor select bus. *McCrory* at col. 7, lines 34-43 utilizes the term "external interrupt;" As stated by Examiner at paragraph 11 of the Action, "interrupt pins would be required to communicate the interrupts to the system."

Examiner's statement clearly indicates that Examiner is referring to a different type of "pin" than the connection pins recited by Appellants' claimed invention. *McCrory* utilizes that term to refer to the standard functionality of an external interrupt, which is inherently very different from a processor connector pin that provides the heterogeneous expansion functionality described in Appellants' claimed invention.

Notably, none of the other referenced sections of *McCrory* even suggest (1) a pin that allows selection of a processor as a master processor when the pin is set to an active state or (2) special buses for connecting a processor and providing additional functionality to the connected processor. Appellants' have clearly defined the specific use of the pins and associated buses as ones having different functionality than a traditional interrupt pin within a processor.

Col. 1, lines 15 to col. 2, line 22 and col. 3 lines 7-17 of *Derrick* also fails to teach or suggest the various buses recited in Appellants' claims, in particular the master processor select bus and extended snoop response bus. A careful reading of these sections of *Derrick* reveals that *Derrick* is totally devoid of any such reference or suggestion.

For the above reasons, Examiner's rejection of Claims 4-5, 11, and 17-18 is not well founded and should be reversed.

iii. Claims 6, 19 (and 7, 20) (R cache line state)

Claim 6 recites:

pin is set when a read operation is issued ...; and snooped by a second added processor with cache line in the R coherency state, ... drives the extended snoop response bus with shared intervention information...
(emphases added).

While *Derrick* provides some description of snooping read operations, there is absolutely no mention or suggestion of setting pins to a particular state and driving an "extended snoop response bus" (not provided by *Derrick's* bus configuration) with shared intervention when the second processor has the cache line in the R state. Examiner attributes to *Derrick* features that

are not taught by nor suggested by *Derrick*. Thus, Examiner's rejection of Claims 6, 7, 19 and 20 is not well founded and should be reversed.

iv. Claims 8 and 16

With respect to Claims 8 and 16, there is no teaching or suggestion of sectoring all caches into widths representing a smallest width cache line within the overall system when heterogeneous processors with associated heterogeneous configuration of caches are later connected to the system. The cited sections of *Derrick* (i.e., col. 1 lines 15-col. 2, line 22 and col. 3, line 7-17) speak generally about caches but fail to suggest this sectoring feature. Thus, Examiner's rejection of Claims 8 and 16 is not well founded and should be reversed.

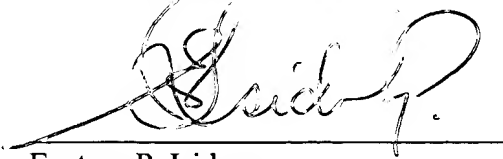
v. Claim 9 and 13 (Switch with director point-to-point processor connections)

As recited by exemplary Claim 9, the processing system comprises "a switch that provides direct point-to-point connection between said first processor and later added processors." The "bus interface system" provided by *McCrory* is a standard multi-drop bus (*see* Figure 3 and 4) that does not provide direct point-to-point connections between individual components, particularly processor-to-processor connections. One skilled in the art is familiar with a switch architecture that provides direct point-to-point connection among components and would not find *McCrory's* standard multi-drop bus to be in any way suggestive of this configuration. Thus, Examiner's rejection of Claims 9 and 13 is not well founded and should be reversed.

CONCLUSION

Appellants have pointed out with specificity the manifest error in the Examiner's rejections, and the claim language that renders the invention patentable over the combination of references. Appellants, therefore, respectfully request that this case be remanded to the Examiner with instructions to issue a Notice of Allowance for all pending claims.

Respectfully submitted,

A handwritten signature in black ink, appearing to read "E. Isidore", written over a horizontal line.

Eustace P. Isidore
Reg. No. 56,104
DILLON & YUDELL LLP
8911 N. Capital of Texas Highway
Suite 2110
Austin, Texas 78759
512-343-6116

ATTORNEY FOR APPELLANTS

APPENDIX

1. A data processing system comprising:

a first processor with a first operational characteristics on a system planar;

interconnection means for later connecting a second processor on said system planar, wherein, when said second processor is heterogenous to said first processor, said interconnection means enables said first processor and said second, heterogenous processor to collectively operate as a symmetric multiprocessor (SMP) system; and

an enhanced operating system (OS) that supports inter-processing operations between said first processor and said second processor including cache coherency operations based on a collective memory configuration of the SMP, wherein said OS logs operating characteristics and cache topology data of each processor connected to the interconnection means to calculate a most efficient work allocation among processors;

wherein said interconnection means and said enhanced operating system support backward and forward compatibility amongst said first processor and said second, heterogenous processor and provides system centric enhancements for inter-processor operations including cache intervention, prefetching, and intelligent cache states.

2. The data processing system of Claim 1, further comprising a second, heterogenous processor connected to said system bus via said interconnect means, wherein said second, heterogenous processor includes different physical component parameters and operational characteristics than said first processor, wherein said different physical component parameters include one or more of a higher number of cache levels, larger cache sizes, improved cache hierarchy, cache intervention, and larger number of on-chip processors.

3. The data processing system of Claim 1, further comprising a cache coherency protocol that supports non-homogenous cache configurations amongst heterogenous processors, said non-homogenous cache configurations including one or more of:

a first cache of the first processor begin designed to support a first set of cache/memory operations with an associated first set of coherency states while a second cache of the second

processor is designed to support a similar set of memory operations with additional coherency states;

said second cache supporting cache intervention from similarly configured caches;

different levels of caches, cache states, and shared caches among processors;

different cache sizes and cache line widths, wherein a first cache line of the first processor's cache having a different width from a cache line of the second processors cache.

4. The data processing system of Claim 3, wherein said interconnect means is coupled to a system bus and comprises a plurality of buses for connecting additional processors to said system bus, said buses comprising system data bus, base address bus, master processor select bus, base snoop response bus, and extended snoop response bus, wherein each bus includes one or more pins that are set/reset to indicate a particular condition of a connected component.

5. The data processing system of Claim 4, wherein said master processor select bus includes a first set of pins, each connected to an added processor, wherein when one of said pins is set to an active state, the connected processor operates as a master on the master processor select bus.

6. The data processing system of Claim 5, wherein:

a respective pin is set when a read operation is issued to indicating that the issuing processor is the master processor; and

when said read operation is snooped by a second added processor with cache line in the R coherency state, the second added processor drives the extended snoop response bus with shared intervention information and sends a retry response on the base snoop response bus.

7. The data processing system of Claim 6, wherein said operational characteristics includes one or more of:

operating frequency, wherein the second processor operates at a higher frequency than said first processor ; and

an instruction ordering mechanism, wherein said first processor and second processor utilizes a different one of a plurality of instruction ordering mechanisms from among in-order processing, out-of-order processing, and robust out-of-order processing.

8. The data processing system of Claim 3, wherein all caches are sectorized into widths representing a smallest width cache line that is accessible within the overall data processing system.

9. The data processing system of Claim 1, further comprising a switch that provides direct point-to-point connection between said first processor and later added processors.

10. A method for upgrading processing capabilities of a data processing system comprising:
providing a plurality of pins from a system bus on a system planar to allow later addition of other processors;

enabling direct connection by a heterogeneous processor to said system planar via said interrupt pins, wherein said interrupt pins provide communication paths between said heterogeneous processor and other processors previously attached to said system planar; and

providing support for full backward compatibility by said new, heterogeneous processor when said new processor comprises more advanced operational characteristics to enable said data processing system to operate as a symmetric multiprocessor system, wherein said support includes an enhanced operating system (OS) that supports inter-processing operations between said first processor and said second processor including cache coherency operations based on a collective memory configuration of the SMP, wherein said OS logs operating characteristics and cache topology data of each processor connected to the interconnection means to calculate a most efficient work allocation among processors;

wherein said interconnection means and said enhanced operating system support backward and forward compatibility amongst said first processor and said second, heterogeneous processor and provides system centric enhancements for inter-processor operations including cache intervention, pre-fetching, and intelligent cache states.

11. The method of Claim 7, wherein said interconnect means is coupled to a system bus and comprises a plurality of buses for connecting additional processors to said system bus, said buses comprising system data bus, base address bus, master processor select bus, base snoop response bus, and extended snoop response bus, wherein each bus includes one or more pins that are set/reset to indicate a particular condition of a connected component;

wherein said master processor select bus includes a first set of pins, each connected to an added processor, wherein one of said pins is set to an active state when the connected processor operates as a master on the master processor select bus;

wherein when said one pin is set when a read operation is issued to indicating that the issuing processor is the master processor; and

when said read operation is snooped by a second added processor in the R coherency state, the second added processor drives the extended snoop response bus with shared intervention information and sends a retry response to on the base snoop response bus.

12. A multiprocessor system comprising:

a plurality of heterogenous processors with different operational characteristics and physical topology connected on a system planar;

a system bus that supports system centric operations;

interrupt pins coupled to said system bus that provide connection for at least one of said plurality of heterogenous processors;

an enhanced system bus protocol that supports downward compatibility of newer processors that designed with advanced operational characteristics from among said plurality of processors to processors that do not support said advance operation characteristics;

an enhanced operating system (OS) that supports inter-processing operations between said first processor and said second processor including cache coherency operations based on a collective memory configuration of the SMP, wherein said OS logs operating characteristics and cache topology data of each processor connected to the interconnection means to calculate a most efficient work allocation among processors;

wherein said enhanced system bus protocol and said enhanced operating system support backward and forward compatibility amongst a first processor and a second, heterogenous processor and provides system centric enhancements for inter-processor operations including cache intervention, prefetching, and intelligent cache states.

13. The multiprocessor system of Claim 12, further comprising a switch that provides direct point-to-point connection between each of said plurality of processors and later added processors.

14. The multiprocessor system of Claim 12, wherein said plurality of processors includes heterogenous processor topologies including different cache sizes, cache states, number of cache levels, and number of processors on a single processor chip.

15. The multiprocessor system of Claim 13, further comprising a cache coherency protocol that supports non-homogenous cache configurations amongst heterogenous processors, said non-homogenous cache configurations including one or more of:

a first cache of the first processor being designed to support a first set of cache/memory operations with an associated first set of coherency states while a second cache of the second processor is designed to support a similar set of memory operations with additional coherency states;

said second cache supporting cache intervention from similarly configured caches;

different levels of caches, cache states, and shared caches among processors;

different cache sizes and widths of cache lines, wherein a first cache line of the first processor's cache having a different width from a cache line of the second processors cache.

16. The multiprocessor system of Claim 15, wherein all caches are sectorized into widths representing a smallest width cache line that is accessible within the overall data processing system.

17. The multiprocessor system of Claim 15, wherein said interconnect means is coupled to a system bus and comprises a plurality of buses for connecting additional processors to said system bus, said buses comprising system data bus, base address bus, master processor select bus, base snoop response bus, and extended snoop response bus, wherein each bus includes one or more pins that are set/reset to indicate a particular condition of a connected component.

18. The multiprocessor system of Claim 17, wherein said master processor select bus includes a first set of pins, each connected to an added processor, wherein when one of said pins is set to an active state, the connected processor operates as a master on the master processor select bus.

19. The multiprocessor system of Claim 18, wherein:

the respective pin is set when a read operation is issued to indicating that the issuing processor is the master processor; and

when said read operation is snooped by a second added processor with cache line in the R coherency state, the second added processor drives the extended snoop response bus with shared intervention information and sends a retry response on the base snoop response bus.

20. The multiprocessor system of Claim 19, wherein said operational characteristics includes one or more of:

operating frequency, wherein the second processor operates at a higher frequency than said first processor; and

an instruction ordering mechanism, wherein said first processor and second processor utilizes a different one of a plurality of instruction ordering mechanisms from among in-order processing, out-of-order processing, and robust out-of-order processing.